

# Capturing exponential variance using polynomial resources: applying tensor networks to non-equilibrium stochastic processes

T. H. Johnson,<sup>1,2,3,4,\*</sup> T. J. Elliott,<sup>2</sup> S. R. Clark,<sup>2,1,3</sup> and D. Jaksch<sup>2,1,3</sup>

<sup>1</sup>Centre for Quantum Technologies, National University of Singapore, 3 Science Drive 2, 117543, Singapore

<sup>2</sup>Clarendon Laboratory, University of Oxford, Parks Road, Oxford OX1 3PU, United Kingdom

<sup>3</sup>Keble College, University of Oxford, Parks Road, Oxford OX1 3PG, United Kingdom

<sup>4</sup>Institute for Scientific Interchange, Via Allassio 11/c, 10126 Torino, Italy

(Dated: April 7, 2015)

Estimating the expected value of an observable appearing in a non-equilibrium stochastic process usually involves sampling. If the observable's variance is high, many samples are required. In contrast, we show that performing the same task without sampling, using tensor network compression, efficiently captures high variances in systems of various geometries and dimensions. We provide examples for which matching the accuracy of our efficient method would require a sample size scaling exponentially with system size. In particular, the high-variance observable  $e^{-\beta W}$ , motivated by Jarzynski's equality, with  $W$  the work done quenching from equilibrium at inverse temperature  $\beta$ , is exactly and efficiently captured by tensor networks.

**Introduction.**—Dynamical stochastic processes are used throughout the natural and social sciences when inaccessible degrees of freedom are well-represented by random variables [1, 2]. To calculate expected observable values, numerical methods are usually required. Out of equilibrium, the typical method is dynamical Monte Carlo [3–8]. Essentially, averaging over randomly sampled paths provides an unbiased estimate for the expected value of an observable. To obtain a fixed expected fractional error, the number of paths sampled must scale linearly with the variance divided by the square of the expected value. For a multitude of important observables, such as those appearing in the estimation of free energies [9] and likelihoods of rare events [10–15], this ratio is large: such observables are said to have high variance and sampling methods struggle when applied to them.

Here we present an approach that is very different to sampling. We simultaneously follow all paths, which is made efficient by using controlled data compression, usually approximate but exact in special cases, based on tensor networks. While tensor networks have previously been used in conjunction with stochastic processes [16–20], the question of how their performance relates to variance has remained unanswered. Understanding this is crucial if we are to know whether or not tensor networks, which have had a revolutionary effect in simulating quantum systems [21–27] and have been used to great effect in solving partial differential equations [28–31], provide a useful and perhaps essential complementary technique to sampling in stochastic processes.

In this Letter we address this question and our answer is very clear: high variance does not limit the accuracy of tensor network compression, and tensor networks can be applied efficiently to tackle problems, even out of equilibrium, for which sampling-based methods struggle. This opens the door for the use of tensor network methods on




FIG. 1. (color online) *Tensor network compression.* (a) An Ising system whose degrees of freedom (blue circles with arrows) interact, in this case, with a two-dimensional lattice geometry (red lines). (b) The probability distributions  $P_E(\mathbf{z})$  and  $P(\mathbf{z}, t)$ , and  $Q(\mathbf{z}, t)$  (see main text) at any time  $t$  are compressed by representing them (approximately or, in special cases, exactly) by a contraction of tensors (green circles) with the same geometry as the interactions. Each black leg corresponds to an index of a tensor, and the joining of two legs represents the contraction of the two corresponding indices.

a wide-variety of non-equilibrium stochastic systems for which capturing high variance is essential. In particular, we show that a distribution of weighted expectation values of high-variance observable  $e^{-\beta W}$ , with  $W$  the work done quenching from equilibrium at inverse temperature  $\beta$ , is represented exactly by a highly compressed tensor network.

We focus on an Ising system, an example of which is shown in Fig. 1(a). It comprises  $N$  nodes, labeled by  $\ell$ , the configuration  $z_\ell$  of each taking one of  $d = 2$  discrete values  $z_\ell = \{-1, 1\}$ . The configuration of all  $N$  nodes is given by the  $N$ -tuple  $\mathbf{z} = (z_1, \dots, z_N)$  and the probability of being in configuration  $\mathbf{z}$  is  $P(\mathbf{z})$ . Tensor networks best suit systems for which crucial quantities, like energy, are  $n$ -bodied, with  $n$  small. We consider the simplest non-trivial case of an energy comprising single and two-body terms

$$E(\mathbf{z}) = -J \sum_{(\ell, \ell') \in \mathcal{E}} z_\ell z_{\ell'} - \lambda \sum_{\ell} z_\ell, \quad (1)$$

where  $\mathcal{E}$  are  $N_{\mathcal{E}}$  edges connecting interacting nodes.

\* tomihjohnson@gmail.com

*Equilibrium.*—The relationship between compressibility and variance out of equilibrium builds on that in equilibrium. The equilibrium Gibbs distribution at inverse temperature  $\beta$  is  $P_E(\mathbf{z}) = e^{-\beta E(\mathbf{z})}$ , normalized to the partition function  $Z_E = \sum_{\mathbf{z}} e^{-\beta E(\mathbf{z})}$ . It is always possible to represent  $P_E(\mathbf{z})$  (or any distribution) by a tensor network of the form shown in Fig. 1(b)

$$P_E(\mathbf{z}) = \sum_{\mathbf{i}} \prod_{\ell} A_{\mathbf{i}_{\ell}}^{[\ell]z_{\ell}}, \quad (2)$$

that shares the same geometry as the interactions. Here  $A^{[\ell]}$  is a tensor associated with node  $\ell$ . It has a physical index  $z_{\ell}$  and  $k_{\ell}$  auxiliary indices  $\mathbf{i}_{\ell}$ , one for each edge connected to  $\ell$ , and each taking one of  $\chi$  values, which may in principle be large. The sum is over the values taken by all auxiliary indices, which is just a sum over an  $N_{\mathcal{E}}$ -tuple  $\mathbf{i}$  of indices. The Gibbs distribution is important because the tensor network representation is exact for  $\chi = d$  (see Supplemental Material [32]). This implies that  $\sum_{\ell} d^{k_{\ell}+1}$  numbers may be used to represent  $d^N$  others, providing a significant yet exact compression if the degrees  $k_{\ell}$  are limited, as in lattice systems with local interactions.

As well as compression, tensor networks offer a means of calculating the partition function  $Z_E$ , since  $Z_E = \sum_{\mathbf{z}} P_E(\mathbf{z}) = \sum_{\mathbf{i}} \prod_{\ell} T_{\mathbf{i}_{\ell}}^{[\ell]}$  with transfer tensors  $T^{[\ell]} = \sum_{z_{\ell}} A^{[\ell]z_{\ell}}$  [33–36]. For a one-dimensional (1D) chain the partition function  $Z_E$  relates to a product of transfer matrices and requires  $\mathcal{O}(Nd^2)$  or  $\mathcal{O}(Nd^3)$  resources to compute for open or periodic boundaries, respectively. In higher dimensions (if the tensor network has a large treewidth [37, 38]) the tensor contractions cannot in general be performed both exactly and efficiently, but efficient strategies exist to perform them approximately. Levin and Nave [35] demonstrated that this can be done accurately for two-dimensional (2D) non-critical lattice systems using tensor renormalization group [39, 40].

Contrastingly, estimating the partition function  $Z_E$  directly by evaluating the sum  $Z_E = \sum_{\mathbf{z}} P_E(\mathbf{z})$  through random sampling is made difficult by the fact that, in general, the variance of observables requiring estimation grows quickly with system size [9]. The tensor network representations of  $P_E(\mathbf{z})$  and  $Z_E$  show that high variance does not imply difficulty in equilibrium, away from criticality.

*Non-equilibrium.*—Out of equilibrium, the dynamics of a Markovian system [41] depends only on its current configuration, and the evolution of the distribution  $P(\mathbf{z}, t)$  is described by a master equation of the form

$$\frac{\partial P(\mathbf{z}, t)}{\partial t} = \sum_{\mathbf{z}'} H(\mathbf{z}, \mathbf{z}', t) P(\mathbf{z}', t). \quad (3)$$

Each non-negative off-diagonal element  $H(\mathbf{z}, \mathbf{z}', t)$  for  $\mathbf{z} \neq \mathbf{z}'$  is the Poisson rate of a transition from  $\mathbf{z}'$  to  $\mathbf{z}$  at time  $t$ , and together these fix the non-positive diagonals  $H(\mathbf{z}, \mathbf{z}, t) = -\sum_{\mathbf{z}' \neq \mathbf{z}} H(\mathbf{z}', \mathbf{z}, t)$  such that the normalization of  $P(\mathbf{z}, t)$  is conserved.  $H$  is commonly referred to as the Hamiltonian.

To simulate such dynamics using non-equilibrium tensor network methods, we represent  $P(\mathbf{z}, t)$  at any time by a tensor network, as in Eq. (2), with a small  $\chi$ . Doing so assumes that this representation, while not necessarily exact, is accurate. There is no guarantee of this accurate compressibility on all occasions, but it is expected in many situations. For example, consider a quench from one Hamiltonian  $H(\mathbf{z}, \mathbf{z}', 0) = H_0(\mathbf{z}, \mathbf{z}')$  to another  $H(\mathbf{z}, \mathbf{z}', \tau) = H_1(\mathbf{z}, \mathbf{z}') \neq H_0(\mathbf{z}, \mathbf{z}')$ , where the system begins in the stationary state  $P_0(\mathbf{z})$  satisfying  $\sum_{\mathbf{z}'} H_0(\mathbf{z}, \mathbf{z}') P_0(\mathbf{z}') = 0$ . For much later times  $t \gg \tau$  (on the timescale  $h^{-1}$ , where  $h$  is some Hamiltonian-specific convergence rate), the system will converge to another stationary state  $P_1(\mathbf{z})$  satisfying  $\sum_{\mathbf{z}'} H_1(\mathbf{z}, \mathbf{z}') P_1(\mathbf{z}') = 0$ . Numerous examples have revealed that stationary states of local stochastic processes are accurately compressible via tensor network representations [16–20]. Thus in such quenches both initial and long-time distributions  $P_0(\mathbf{z})$  and  $P_1(\mathbf{z})$  are accurately compressible. Unlike for quantum systems [42], compression errors are limited even when a system is driven away from equilibrium. The example on which we focus here is that of a thermalizing (equilibrating) Hamiltonian  $H(\mathbf{z}, \mathbf{z}', t) = H_E(\mathbf{z}, \mathbf{z}')$  for which the Gibbs distribution is stationary,  $\sum_{\mathbf{z}'} H_E(\mathbf{z}, \mathbf{z}') P_E(\mathbf{z}') = 0$ , and it is the energy  $E(t)$  that is quenched by varying the bias  $\lambda(t)$  appearing in Eq. (1).

The probability distribution  $P(\mathbf{z}, t)$  over configurations contains only partial information about the full probability distribution over the possible paths through configuration space taken by the stochastic process. As such, the expected values of only certain observables may be calculated from  $P(\mathbf{z}, t)$ . These include observables whose values  $O(\mathbf{z})$  depend on the configuration  $\mathbf{z}$  of the system at a single time  $t$ , thus having expected value  $\langle O(t) \rangle = \sum_{\mathbf{z}} O(\mathbf{z}) P(\mathbf{z}, t)$ . We call such observables configuration-dependent and use the example of the magnetization  $M(\mathbf{z}) = \sum_{\ell} z_{\ell}$ . The values of some other observables depend on the full path taken by the system and their expected values cannot be calculated from  $P(\mathbf{z}, t)$ . We call such observables path-dependent, and use the example of the work done  $W(t) = -\int_0^t ds M(\mathbf{z}(s)) \dot{\lambda}(s)$  by varying  $\lambda(t)$  between times 0 and  $t$ .

Although not previously considered in the literature, the expected values of some path-dependent observables can indeed be calculated using tensor networks and, as we will show, provide us with a stark example of exact compressibility in the face of high variance out of equilibrium. The idea is to represent the relevant path-dependent information locally in time, not with  $P(\mathbf{z}, t)$ , but through the distribution of weighted conditional expected values  $Q(\mathbf{z}, t) = P(\mathbf{z}, t) \langle O(\mathbf{z}, t) \rangle$ . Here  $\langle O(\mathbf{z}, t) \rangle$  is the expected value an observable has accumulated by time  $t$  conditioned upon the system arriving at configuration  $\mathbf{z}$  at that time. The expected value of interest  $\langle O(t) \rangle = \sum_{\mathbf{z}} Q(\mathbf{z}, t)$  is then obtained from  $Q(\mathbf{z}, t)$ . It follows from Eq. (3) that the distribution  $Q(\mathbf{z}, t)$  evolves

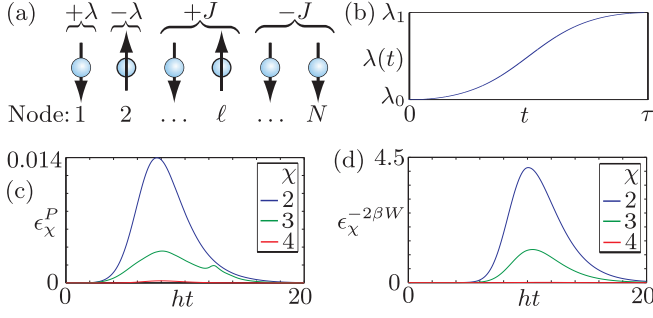


FIG. 2. (color online) *Accurate compression out of equilibrium.* (a) An  $N$ -node open Ising chain with exchange  $J$  and bias  $\lambda$ . (b) The bias  $\lambda(t)$  is varied, driving the system away from equilibrium. (c) The fractional error  $\epsilon_1$  between calculating  $\langle e^{M(t)} \rangle$  using probability distribution  $P(\mathbf{z}, t)$  and its compressed tensor network approximation of dimension  $\chi$  (see Supplemental Material [32]). A dashed line marks the end of the quench at time  $t = \tau$ . (d) Similarly, the fractional error  $\epsilon_2$  between calculating  $\langle e^{-2\beta W(t)} \rangle$  by summing  $Q(\mathbf{z}, t)$  and its compressed tensor network approximation. The parameters used are  $\beta\lambda_0 = 0$ ,  $\beta\lambda_1 = 1$ ,  $h\tau = 10$ ,  $\beta J = 1$  and  $N = 8$ .

as

$$\frac{\partial Q(\mathbf{z}, t)}{\partial t} = \sum_{\mathbf{z}'} H'(\mathbf{z}, \mathbf{z}', t) Q(\mathbf{z}', t), \quad (4)$$

with  $H'(\mathbf{z}, \mathbf{z}', t) = H(\mathbf{z}, \mathbf{z}', t) + \dot{o}(\mathbf{z}, t)\delta(\mathbf{z}, \mathbf{z}')$ , where  $\dot{o}(\mathbf{z}, t)$  is the rate of increase of the natural logarithm of the observable at configuration  $\mathbf{z}$  and time  $t$ .

It is desirable to predict, as we have for  $P(\mathbf{z}, t)$ , the accuracy of compressing  $Q(\mathbf{z}, t)$  at any time during its evolution using tensor networks. Consider the quench in  $\lambda(t)$  between times 0 and  $\tau$ , starting from equilibrium. The distributions are initially equal  $Q(\mathbf{z}, 0) = P(\mathbf{z}, 0)$ , thus the accurate compressibility of the latter implies the same of the former. Additionally,  $\dot{o}(\mathbf{z}, t)$  is only non-zero for times  $t < \tau$  and the stochastic evolution is ergodic. Thus after a sufficiently long time  $t \gg \tau$  (relative again to convergence timescale  $h^{-1}$ ) the configurations will have mixed such that  $\langle O(\mathbf{z}, t) \rangle = \langle O(t) \rangle$  is independent of  $\mathbf{z}$  and thus once again  $Q(\mathbf{z}, t) = \langle O(t) \rangle P(\mathbf{z}, t)$  is as accurately compressible as  $P(\mathbf{z}, t)$ .

*Numerical examples.*—We demonstrate these behaviors for a system undergoing thermalizing Glauber dynamics [43] via local transitions,

$$H_E(\mathbf{z}, \mathbf{z}') = h(\mathbf{z}, \mathbf{z}') \left[ 1 + e^{-\beta(E(\mathbf{z}') - E(\mathbf{z}))} \right]^{-1},$$

for  $\mathbf{z} \neq \mathbf{z}'$ . Here  $h(\mathbf{z}, \mathbf{z}') = h(\mathbf{z}', \mathbf{z})$  are symmetric rates equaling a non-zero rate  $h$  only where  $\mathbf{z}$  and  $\mathbf{z}'$  differ by the configuration of a single node. The energy is quenched via the parameter  $\lambda(t)$  varying from  $\lambda_0$  to  $\lambda_1$  over time  $0 \leq t \leq \tau$  according to a smoothed tanh ramp (see Supplemental Material [32]), as drawn in Fig. 2(b). We focus on configuration-dependent observable  $e^{M(t)}$  and path-dependent observables  $e^{-2\beta W(t)}$  and  $e^{-\beta W(t)}$ .

All have variance over mean squared growing exponentially with system size  $N$ . Initially, we consider the Ising nodes to be in an open 1D chain, illustrated in Fig. 2(a).

To assess the accuracy of compression, we exactly calculate distributions  $P(\mathbf{z}, t)$  and  $Q(\mathbf{z}, t)$  for small  $N = 8$  at time  $t$ , then calculate the error in the expected value of observables induced by compressing the distributions as a tensor network. The errors shown in Figs. 2(c) and (d) for  $e^{M(t)}$  and  $e^{-2\beta W(t)}$ , respectively, show that, despite large variances, expected values are relatively unaffected by tensor network compression. The distributions are exactly compressible at  $t = 0$  and thus no error occurs, as expected. The errors due compression initially increase as  $\lambda(t)$  varies, then decrease exponentially to small values again on a timescale  $\sim h^{-1}$ . Interestingly, errors begin to decrease even at times  $t < \tau$  prior to the end of the quench. For  $\chi \gtrsim 4$  the compression is near-exact at all times. We arrive at similar conclusions for other types of variation tried e.g. linear, variations in  $J$ , and other observables.

A striking example is found in the path-dependent observable  $e^{-\beta W(t)}$ . The observable has received particular attention due to its featuring in several non-equilibrium identities in statistical physics, such as that by Jarzynski [44]. Crucially for our discussion, the *non-equilibrium* distribution  $Q(\mathbf{z}, t) = P(\mathbf{z}, t)\langle O(\mathbf{z}, t) \rangle$  for this special case has an *equilibrium* structure  $Q(\mathbf{z}, t) = P_t(\mathbf{z})Z_t(\beta)/Z_0(\beta)$  [44], where we have used shorthands of the form  $P_t(\mathbf{z})$  for the Gibbs distribution corresponding to  $\lambda(t)$  and  $Z_t(\beta)$  for the corresponding partition function (where from now on we normalize the Gibbs distribution to 1). It immediately follows from our discussion of systems in equilibrium that  $Q(\mathbf{z}, t)$ , despite containing information about non-equilibrium high-variance observables, has an exact highly-compressed  $\chi = d$  tensor network representation at all times  $t$ . This can be used to efficiently and accurately calculate not only  $\langle O(t) \rangle$  but a range of properties of the work distribution during such dynamics. Note that this exact behavior is particular to  $e^{-\beta W(t)}$  and doesn't even extend to its square  $e^{-2\beta W(t)}$ .

We have so far demonstrated accurate single-time compressibility. We next examine how this extends to a dynamical tensor network simulation, where compression of  $P(\mathbf{z}, t)$  or  $Q(\mathbf{z}, t)$  occurs not only at a single time but at all times during their evolution. While one might expect the compression errors at single times to accumulate, we find this is mitigated by the ergodicity of the evolution (unlike in quantum systems). For example, errors in  $P(\mathbf{z}, t)$  will not change the distribution to which it converges, and thus the significance of transient errors diminish, rather than accumulate, in time. In what follows, the specific algorithm we use to perform the evolutions of Eqs. (3) and (4) is time-evolution block decimation (TEBD) [20, 45, 46]. The TEBD algorithm uses a timestep  $\delta t$  resulting in an error, beyond that due to compression, of  $\mathcal{O}(N\delta t^2)$  and requires time  $\mathcal{O}(N\chi^3\delta t^{-1})$  (see the Supplemental Material [32]).

We first calculated  $\langle e^M \rangle$  and  $\langle e^{-2\beta W} \rangle$ , where  $M =$






FIG. 3. (color online) *Compression and variance.* (a) The fractional error  $\epsilon_1$  of the calculated expected value ( $e^M$ ) relative to the  $\chi = 10$  result. (b) The analogous error  $\epsilon_2$  for observable  $e^{-2\beta W}$ . (c) The variance over mean squared  $v_3$  of  $e^{-\beta W}$  as a function of  $h\tau$  and  $N$ , obtained by calculating  $\langle e^{-2\beta W} \rangle$  using  $\chi = 10$  and  $\delta t = 10^{-2}$ , and using an exact result for  $\langle e^{-\beta W} \rangle$ . (d) For a periodic  $N = 64 \times 64$  system, the expectation value  $\langle \exp^{-\beta W(t)} \rangle_S$ , conditional on one spin taking value  $z_1 = 1$  at  $t$  (full line). Also, the ratio  $r_S(t)$  of this value relative to that for a reversible quench (dashed line). Unless stated otherwise the parameters used are  $\beta\lambda_0 = 0$ ,  $\beta\lambda_1 = 1$ ,  $h\tau = 10$ ,  $\beta J = 1$ ,  $N = 200$  and  $h\delta t = 10^{-3}$ .

$M(\tau)$  and  $W = M(\tau)$  are values at the end of the quench  $t = \tau$ . There are no exact values available to compare against for large systems, but the compressibility expected from Figs. 2(c) and (d) is confirmed by our TEBD results in Figs. 3(a) and (b), respectively. The calculated expected values converge approximately exponentially with increasing  $\chi$ , reaching acceptably converged values by  $\chi \lesssim 5$ , despite the variance over mean squared being very large for the  $N = 200$  considered.

We next calculated  $\langle e^{-\beta W} \rangle$ . Since compression is exact for this observable if  $\chi \geq 2$ , the error  $\epsilon_3$  is purely due to the finite timestep  $\delta t$  and, as stated above, scales as  $\epsilon_3 \propto N$ . Meanwhile the estimated variance over mean squared  $v_3$  scales with  $e^{O(N)}$ , as shown in Fig. 3(c). It is then particularly clear with this example that, to achieve the same fractional accuracy (namely,  $\epsilon_3 \propto N$ ) as we efficiently obtain, a naïve sampling method would require a sample size and thus time  $e^{O(N)}$  in contrast to the TEBD algorithm that requires time  $O(N)$ . Explicitly, our  $N = 200$ ,  $\chi = 2$ ,  $h\delta t = 10^{-4}$ ,  $h\tau = 10$  calculation takes less than an hour and achieves an error  $\epsilon_3 \approx 10^{-6}$ . We do not compare this against the cpu time of any one sampling algorithm, as this choice is likely to be unrepresentative. Instead, we note that, since  $v_3 \approx 10^{11}$ , our accuracy would require  $\approx 10^{17}$  samples to reproduce, and so matching our cpu time would require each path to be sampled in  $\approx 10^{-13}$ s.

As our final example, we demonstrate both the diversity of information stored in the distribution  $Q(\mathbf{z}, t)$  and the application of our method to 2D lattice geometries,

specifically a  $N = 64 \times 64$  square periodic lattice. We consider  $\langle O(t) \rangle_S = \sum_{\mathbf{z} \in S} P(\mathbf{z}, t) \langle O(\mathbf{z}, t) \rangle / \sum_{\mathbf{z} \in S} P(\mathbf{z}, t)$ , the expected value of  $O(t) = e^{-\beta W(t)}$  given the system's configuration  $\mathbf{z}$  is in subset  $S$  at time  $t$ , where  $S$  is the set configurations in which one node has value  $z_1 = 1$ . The observable has a very large variance. Its conditional expected value can be rewritten  $\langle O(t) \rangle_S = \sum_{\mathbf{z} \in S} Q(\mathbf{z}, t) / \sum_{\mathbf{z} \in S} P(\mathbf{z}, t)$ . The numerator  $\sum_{\mathbf{z} \in S} Q(\mathbf{z}, t)$  corresponds to a high-variance observable but can nevertheless be efficiently evaluated using equilibrium techniques for Gibbs distributions, due to the equilibrium structure of the non-equilibrium distribution  $Q(\mathbf{z}, t)$ . In this case we use the tensor renormalization group method [35] with intermediary dimension  $\chi = 3$  (see Supplemental Material [32]), which suffices as the model parameters are far from criticality. The denominator  $\sum_{\mathbf{z} \in S} P(\mathbf{z}, t)$  is simply the low-variance expected value of an observable taking value unity when  $\mathbf{z}(t) \in S$ , otherwise zero. This can be accurately calculated using our tensor network methods or the more common dynamical Monte Carlo methods [5–7]. We used the latter, with sample size  $\sim 1.5 \times 10^4$  (see Supplemental Material [32]). The result is plotted in Fig. 3(d). Also plotted is  $r_S(t)$ , the ratio of the expected value  $\langle O(t) \rangle_S$  to its value were the system in equilibrium at all times. This emphasizes that, despite exploiting an equilibrium structure, the dynamics being simulated is truly irreversible.

*Discussion.*—We have shown that tensor networks provide a way to overcome the challenges faced by sampling methods when estimating expected values of high-variance observables out of equilibrium, even finding an exactly compressible non-equilibrium example. While advanced techniques for variance reduction exist, such as sequential importance sampling [10–13] and branching methods [14, 15], using these techniques usually requires judicious choices specific to the models to be simulated based on prior intuition about the process. No such intuition or choices are needed for a tensor network calculation. However, whether tensor networks remain accurate and efficient for geometries beyond 1D and 2D lattices, and other models e.g. describing frustration or disorder, must still be established.

Finally let us comment on how our findings relate to the wider use of tensor networks. While outstanding efficiency is possible using dynamical tensor network algorithms for finite 1D pure *quantum* systems, there is an ongoing effort from the community to reach larger dimensions  $\chi$  and sizes  $N$  in 2D. Reference [47] gives a state-of-the-art demonstration in which the ground state of an  $N = 21 \times 21$  system with  $d = 2$  is calculated using  $\chi = 8$ . Meanwhile we have seen here that often very small dimensions  $\chi$  are required to simulate *classical* systems, even during real time dynamics. Further, since only one copy of  $P(\mathbf{z})$  is needed in classical algorithms, compared to two copies of the wavefunction in quantum algorithms, those that take time e.g.  $O(\chi^6)$  for quantum systems will instead take time e.g.  $O(\chi^3)$  for classical systems. It may therefore be the case that clas-

sical stochastic systems are currently in an even better position than quantum systems to benefit from current high-dimensional tensor network algorithms.

THJ, SRC, and DJ thank the National Research Foundation and the Ministry of Education of Singapore for support. TJE thanks St Anne's College of the University of Oxford for financial support. This research received funding from the ERC under the European Unions Seventh Framework Programme (FP7/2007-2013)/ERC Grant Agreement no. 319286 Q-MAC and from the EPSRC through projects EP/K038311/1 and EP/J010529/1.

## SUPPLEMENTAL MATERIAL

### 1. Exact tensor network for a Gibbs distribution

It is possible to represent the Gibbs distribution exactly by a tensor network

$$P_E(\mathbf{z}) = \sum_{\mathbf{i}} \prod_{\ell} A_{\mathbf{i}_{\ell}}^{[\ell]z_{\ell}},$$

where the bond dimensions are  $\chi = d$ .

The generality of such a representation follows from the arguments justifying the generality of the tensor network factorization of the partition function as, for example, appears in Ref. [35]. Here we give explicit expressions for the tensors  $A^{[\ell]}$  corresponding to the two systems discussed in the main text: the one-dimensional (1D) open Ising chain and the two-dimensional (2D) periodic square Ising system, both with  $d = 2$ .

For the 1D open Ising chain the elements of tensors  $A^{[\ell]}$  are

$$\begin{aligned} A_{i_1}^{[1]z_1} &= \exp[\beta z_1 (J i_1 / 2 + \lambda)], \\ A_{i_{\ell-1} i_{\ell}}^{[\ell]z_{\ell}} &= \exp[\beta z_{\ell} (J(i_{\ell-1} + i_{\ell}) / 2 + \lambda)], \quad 1 < \ell < N, \\ A_{i_{N-1}}^{[N]z_N} &= \exp[\beta z_N (J i_{N-1} / 2 + \lambda)]. \end{aligned}$$

The labeling of the indices  $i_{\ell}$ , for  $\ell = 1, \dots, N-1$ , makes it clear which pairs are contracted. The indices take values  $\pm 1$  and so the matrices  $A^{[\ell]z_{\ell}}$  for  $1 < \ell < N$  are  $2 \times 2$  (the boundary vectors  $A^{[1]z_1}$  and  $A^{[N]z_N}$  are  $1 \times 2$  and  $2 \times 1$ , respectively).

For the 2D periodic square Ising system the tensors  $A^{[\ell]}$  are identical and their elements are given by

$$A_{i_1 i_2 i_3 i_4}^{[\ell]z_{\ell}} = \exp \left[ \beta z_{\ell} \left( J \sum_{\ell'=1}^4 i_{\ell'} / 2 + \lambda \right) \right].$$

The indices should be paired to form a square periodic lattice. The order of the indices at any one site is not important since the tensors are symmetric in the indices  $i_1, \dots, i_4$ . As before, the indices  $i_{\ell'}$  for  $\ell' = 1, \dots, 4$  take values  $\pm 1$ .

### 2. Time-dependence of the bias

In our calculations the bias  $\lambda(t)$  is varied between  $\lambda_0$  and  $\lambda_1$  according to a smoothed tanh ramp. The exact time-dependence used is

$$\lambda(t) = \lambda_0 + \frac{\lambda_1 - \lambda_0}{2} \left[ 1 + \frac{\tanh \left( \sin \left[ \left( \frac{t}{\tau} - \frac{1}{2} \right) \pi \right] \right)}{\tanh(1)} \right],$$

where  $\tau$  is the duration of the quench.

### 3. Single-time compression

With the system out of equilibrium, we wish to determine the effect of compressing a distribution  $P(\mathbf{z}, t)$  (similarly  $Q(\mathbf{z}, t)$ ) at some time  $t$ . In other words we wish to construct a tensor network of small dimension  $\chi$  that closely approximates the distribution

$$P(\mathbf{z}, t) \approx \sum_{\mathbf{i}} \prod_{\ell} A_{\mathbf{i}_{\ell}}^{[\ell]z_{\ell}}.$$

For a 1D system with open boundaries the tensor network takes the matrix product form

$$\sum_{\mathbf{i}} \prod_{\ell} A_{\mathbf{i}_{\ell}}^{[\ell]z_{\ell}} = A^{[1]z_1} A^{[2]z_2} \dots A^{[N]z_N},$$

with  $\chi \times \chi$  matrices  $A^{[\ell]z_{\ell}}$  for  $1 < \ell < N$ , and  $1 \times \chi$  and  $\chi \times 1$  boundary vectors  $A^{[1]z_1}$  and  $A^{[N]z_N}$ . It is well-known how to find an accurate matrix product approximation to a distribution  $P(\mathbf{z}, t)$  using singular value decomposition decimation, as explained in Refs. [20, 45]. Note that the matrix product obtained may not be optimally accurate for the observable of interest, and thus even more effective compression may be possible than that we have reported. However, singular value decomposition decimation is closely related to how compression occurs in the operation of the time-evolution block decimation (TEBD) algorithm, and thus we expect the conclusions regarding compressibility at a single time here to be most relevant to repeated compression within the TEBD algorithm.

### 4. Time-evolution block decimation

TEBD is a standard method, detailed in Refs. [20, 45], for evolving a matrix product representation of some distribution  $P(\mathbf{z}, t)$  according to an equation of the form

$$\frac{\partial P(\mathbf{z}, t)}{\partial t} = \sum_{\mathbf{z}'} H(\mathbf{z}, \mathbf{z}', t) P(\mathbf{z}', t).$$

The algorithm proceeds by breaking the evolution into timesteps  $\delta t$ . Evolving the distribution  $P(\mathbf{z}, t)$  over a timestep takes it away from one expressed in terms of

a matrix product of dimension  $\chi$ . Thus the distribution is repeatedly, within each timestep, re-compressed to this form using the singular value decomposition. The method we use to discretize time is a typical one and is explained in detail in Ref. [20]. It is accurate up to errors that are second-order in  $\delta t$ .

The most typical implementations of TEBD are compatible with two-site nearest-neighbor Hamiltonians only. Since the values of the transition rates appearing in Glauber dynamics depend on the energies at three nodes, the Hamiltonian is not naturally two-site. To bring it into standard form, rather than  $N$  single nodes with  $d = 2$ , the system is broken down into  $N/2$  supernodes with  $d = 4$  according to  $\mathbf{z} = (z_{1,2}, \dots, z_{N-1,N})$ , where  $z_{2\ell-1,2\ell} = \{(-1, -1), (-1, 1), (1, -1), (1, 1)\}$  for  $\ell = 1, \dots, N/2$ . In the supernode picture the Hamiltonian is naturally formed of two-site nearest-neighbor terms.

Finally, the calculation results and scaling of resources quoted in the main text is for fixed IEEE 64-bit double precision computing. Though not necessary for the calculations presented here, higher than normal precision may be required for very large systems (see e.g. Ref. [17]) due to the number of orders of magnitude spanned by the observable and probability distributions.

## 5. Tensor renormalization group

We have used the tensor renormalization group method to calculate the sum  $\sum_{\mathbf{z} \in S} Q(\mathbf{z}, t)$  for the case of an  $N = 64 \times 64$  2D periodic square Ising system,  $O(t) = e^{-\beta W(t)}$  and  $S$  the set of configurations with  $z_1 = 1$ . We are able to use tensor renormalization group methods because of the equilibrium structure  $Q(\mathbf{z}, t) = P_t(\mathbf{z})Z_t(\beta)/Z_0(\beta)$  taken by the non-equilibrium distribution  $Q(\mathbf{z}, t)$  for this observable.

To see this, we use the above to rewrite  $\sum_{\mathbf{z} \in S} Q(\mathbf{z}, t) = (Z_t(\beta)/Z_0(\beta)) \sum_{\mathbf{z} \in S} P_t(\mathbf{z})$ , where  $P_t(\mathbf{z})$  is the equilibrium Gibbs distribution corresponding to the bias  $\lambda(t)$  at time  $t$ , and  $Z_t(\beta)$  the corresponding partition function. Thus the calculation reduces to calculating partition functions  $Z_t(\beta)$  and partial sums  $\sum_{\mathbf{z} \in S} P_t(\mathbf{z})$  over Gibbs distributions conditioned upon the value of a single spin  $z_1 = 1$ . Given the representation of a Gibbs distribution  $P_t(\mathbf{z})$  by a tensor network, the partition function

$Z_t = \sum_{\mathbf{z}} P_t(\mathbf{z}) = \sum_{\mathbf{i}} \prod_{\ell} T_{\mathbf{i}_{\ell}}^{[\ell]}$  is merely a contraction of transfer tensors  $T^{[\ell]} = \sum_{z_{\ell}} A^{[\ell]z_{\ell}}$  with a square geometry. The partial sum  $\sum_{\mathbf{z} \in S} P_t(\mathbf{z}) = \sum_{\mathbf{i}} \prod_{\ell} T_{\mathbf{i}_{\ell}}^{[\ell]}$  is similar except with  $T^{[1]} = A^{[\ell]z_1}$ .

These contractions can be performed using the well-known tensor renormalization group method [35], the relevant details about which we provide here. The strategy is first to perform the contractions involving each set of four nodes making up a node of a coarse-grained square lattice. We are then left with tensor network in a square lattice with the number of nodes reduced by a factor of four, but the dimension of each index potentially increased by a power of 4. Next, we compress this to obtain smaller index dimension  $\chi$ . The contraction/compression process is then repeated iteratively until a single node remains, and the contraction can be trivially evaluated.

For  $N = 64 \times 64$  this requires six iterations. We found it acceptable to use an intermediary dimension  $\chi = 3$ .

## 6. Dynamical Monte Carlo

We have used dynamical Monte Carlo to calculate the sum  $\sum_{\mathbf{z} \in S} P(\mathbf{z}, t)$  for the case of an  $N = 64 \times 64$  periodic square Ising system and  $S$  the set of configurations with  $z_1 = 1$ . We were able to use dynamical Monte Carlo because, by defining an observable  $O(\mathbf{z})$  taking values 1 if  $\mathbf{z} \in S$  and 0 otherwise, we have  $\sum_{\mathbf{z} \in S} P(\mathbf{z}, t) = \sum_{\mathbf{z}} O(\mathbf{z})P(\mathbf{z}, t)$  i.e. the quantity of interest is just the expected value  $\langle O(t) \rangle$  of this observable.

We implement a variable-timestep dynamical Monte Carlo algorithm (as e.g. explained in Ref. [7]) to sample  $N_s$  paths and estimate  $\langle O(t) \rangle$  by averaging the values taken by the observable over such paths. Because the variance of the observable is necessarily less than unity, the expected error due to insufficient sampling must be less than  $1/\sqrt{N_s}$ , which is small for the  $N_s = 14790$  we used.

The only source of error other than insufficient sampling is an approximation made when estimating the time until the next transition as part of the sampling of a path. We assumed for simplicity that the transition rates stayed constant between transitions, which ignores their varying due to the change in  $\lambda(t)$ . Since the average time between transitions is on the order of  $(64^2 h)^{-1}$  and  $\beta\lambda$  is varied by unity over a time  $10h^{-1}$  in our example, the approximation induced by this is negligible.

- 
- [1] N. G. Van Kampen, *Stochastic Processes in Physics and Chemistry* (Elsevier, Amsterdam, 2007), 3rd Ed.
  - [2] T. Rolski, H. Schmidli, V. Schmidt, and J. Teugels, *Stochastic Processes for Insurance and Finance* (John Wiley & Sons, Chichester, 1999).
  - [3] D. R. Cox and H. D. Miller, *The Theory of Stochastic Processes* (Methuen, London, 1965), p. 6.

- [4] W. M. Young and E. W. Elcock, Proc. Phys. Soc. **89**, 735 (1966).
- [5] A. B. Bortz, M. H. Kalos, and J. L. Lebowitz, J. Comp. Phys. **17**, 10 (1975).
- [6] D. T. Gillespie, J. Comput. Phys. **22**, 403 (1976).
- [7] A. P. J. Jansen, preprint arXiv:cond-mat/0303028 (2003).

- [8] A. Chatterjee and D. G. Vlachos, J. Comput. Aided Mater. Des. **14**, 253 (2007).
- [9] G. Potamianos and J. Goutsias, IEEE Trans. Inf. Theory **43**, 1948 (1997).
- [10] C. Dellago, P. G. Bolhuis, F. S. Csajka, and D. Chandler, J. Chem. Phys. **108**, 1964 (1998).
- [11] O. Mazonka, C. Jarzynski, and J. Blocki, Nuc. Phys. A **641**, 335 (1998).
- [12] H. K. Lee and Y. Okabe, Phys. Rev. E **71**, 015102(R) (2005).
- [13] A. Kundu, S. Sabhapandit, and A. Dhar, Phys. Rev. E **83**, 031119 (2011).
- [14] P. Grassberger, Comp. Phys. Comm. **147**, 64 (2002).
- [15] T. Dean and P. Dupuis, Stochastic Process. Appl., **119**, 562 (2009).
- [16] Y. Hieida, J. Phys. Soc. Jpn. **67**, 369 (1998).
- [17] E. Carlon, M. Henkel and U. Schollwöck, Eur. Phys. J. B **12**, 99 (1999).
- [18] E. Carlon, M. Henkel and U. Schollwöck, Phys. Rev. E **63**, 036101 (2001).
- [19] K. Temme and F. Verstraete, Phys. Rev. Lett. **104**, 210502 (2010).
- [20] T. H. Johnson, S. R. Clark, and D. Jaksch, Phys. Rev. E **82**, 036702 (2010).
- [21] F. Verstraete, V. Murg and J. I. Cirac, Adv. Phys. **57**, 143 (2008).
- [22] J. I. Cirac and F. Verstraete, J. Phys. A **42**, 504004 (2009).
- [23] U. Schollwöck, Annals of Physics **326**, 96–192 (2011).
- [24] S. Al-Assam, S. R. Clark, C. J. Foot, and D. Jaksch, Phys. Rev. B **84**, 205108 (2011).
- [25] S. J. Denny, J. D. Biamonte, D. Jaksch, and S. Clark, J. Phys. A: Math. Theor. **45**, 015309 (2011).
- [26] E. Evenbly and G. Vidal. *Quantum criticality with the multi-scale entanglement renormalization ansatz*. In E. Avella and F. Mancini (Eds.) *Strongly Correlated Systems: Numerical Methods*, pp. 99–130 (Springer, Heidelberg, 2013).
- [27] R. Orus, *A practical introduction to tensor networks: Matrix product states and projected entangled pair states* arXiv:1306.2164.
- [28] S. Iblisdir, R. Orus, and J.I. Latorre, Phys. Rev. B **75**, 104305, (2007).
- [29] O. Koch *et al.*, SIAM J. Matrix Anal. Appl. **29**, 434 (2007); *ibid.* **31**, 2360 (2010).
- [30] S. Holtz *et al.*, Num. Math. **120**, 701 (2012).
- [31] C. Lubich *et al.*, SIAM J. Matrix Anal. Appl. **34**, 470-494 (2013).
- [32] See Supplemental Material submitted alongside this manuscript.
- [33] T. Nishino, J. Phys. Soc. Jpn. **64**, 3598 (1995); T. Nishino and K. Okunishi, in *Transfer-Matrix Approach to Classical Systems*, edited by I. Peschel, X. Wang, and K. Hallberg, Lecture Note in Physics Vol. **528** (Springer, Berlin, 1999), pp. 127–148.
- [34] T. Nishino and K. Okunishi, J. Phys. Soc. Jpn. **65**, 891 (1996); *ibid.* **66**, 3040 (1997); T. Nishino, K. Okunishi, and M. Kikuchi, Phys. Lett. A **213**, 69 (1996).
- [35] M. Levin and C. P. Nave, Phys. Rev. Lett. **99**, 120601 (2007).
- [36] R. Orús and G. Vidal, Phys. Rev. B **78**, 155117 (2008).
- [37] Y.-Y. Shi, L.-M. Duan, and G. Vidal, Phys. Rev. A **74**, 022320 (2006).
- [38] I. L. Markov and Y.-Y. Shi, SIAM J. Comput. **38**, 963–981 (2008).
- [39] Z. Y. Xie, H. C. Jiang, Q. N. Chen, Z. Y. Weng, and T. Xiang, Phys. Rev. Lett. **103**, 160601 (2009).
- [40] H. H. Zhao, Z. Y. Xie, Q. N. Chen, Z. C. Wei, J. W. Cai, and T. Xiang, Phys. Rev. B **81**, 174411 (2010).
- [41] J. R. Norris, *Markov Chains* (Cambridge University Press, Cambridge, 1998).
- [42] P. Calabrese and J. Cardy, J. Stat. Mech P04010 (2005).
- [43] R. J. Glauber, J. Math. Phys. **4**, 294 (1963).
- [44] C. Jarzynski, Phys. Rev. E **56**, 5018 (1997).
- [45] G. Vidal, Phys. Rev. Lett. **93**, 040502 (2004).
- [46] Tensor Network Theory Library, Beta Version 1.0.0 (2014), S. Al-Assam, S. R. Clark, D. Jaksch and TNT Development team, [www.tensornetworktheory.org](http://www.tensornetworktheory.org).
- [47] M. Lubasch, I. J. Cirac, and M.-C. Bañuls, preprint arXiv:1405.3259 (2014).
- [48] Longer range but still local interaction could also be treated using tensor networks.
- [49] The values of  $A^{[\ell]}$  follow from similar arguments to those for the tensor network factorization of the partition function [35].
- [50] The exact time dependence used is
- $$\lambda(t) = \lambda_0 + \frac{\lambda_1 - \lambda_0}{2} \left[ 1 + \frac{\tanh \left( \sin \left[ \left( \frac{t-t_0}{t_1-t_0} - \frac{1}{2} \right) \pi \right] \right)}{\tanh(1)} \right].$$
- [51] This can be extended easily to include observables whose values change multiplicatively when a transition occurs.
- [52] To allow us to use nearest-neighbor gates only, rather than  $N$  single spins with  $d = 2$ , the system is broken down into  $N/2$  superspins with  $d = 4$  according to  $\mathbf{z} = (z_{1,2}, \dots, z_{N-1,N})$ , where  $z_{2\ell-1,2\ell} = \{(-1, -1), (-1, 1), (1, -1), (1, 1)\}$  for  $\ell = 1, \dots, N/2$ .
- [53] The calculation results and scaling of resources quoted here is for fixed IEEE 64-bit double precision computing. Though not necessary for the calculations presented here, higher than normal precision may be required for very large systems due to the number of orders of magnitude spanned by the observable and probability distributions.